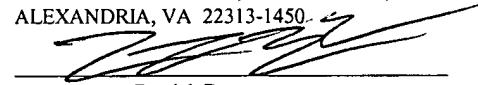


PATENT
5760-20000
VRTS 0629

"EXPRESS MAIL" MAILING LABEL NUMBER
FV318248706 US

DATE OF DEPOSIT 3-2-04

I HEREBY CERTIFY THAT THIS PAPER OR
FEE IS BEING DEPOSITED WITH THE
UNITED STATES POSTAL SERVICE
"EXPRESS MAIL POST OFFICE TO
ADDRESSEE" SERVICE UNDER 37 C.F.R.
§1.10 ON THE DATE INDICATED ABOVE
AND IS ADDRESSED TO THE
COMMISSIONER FOR PATENTS, MAILSTOP
PATENT APPLICATION, P.O. BOX 1450,
ALEXANDRIA, VA 22313-1450.



Derrick Brown

Central Administration and Maintenance of Workstations Using Virtual
Machines, Network Filesystems, and Replication

By:

Hans F. van Rietschote
Mahesh P. Saptarshi
Craig W. Hobbs

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention is related to the field of administering and maintaining computer systems on a network or sometimes connected to a network.

2. Description of the Related Art

An important figure for computer systems is the total cost of ownership (TCO). The TCO includes the costs in acquiring the computer system and related software, and also includes a variety of ongoing costs for maintaining the computer system in working order. Exemplary TCO costs may include costs for: installing the operating system (OS) and the applications needed by the user of the computer system; configuring the applications to meet the user's needs; updating and upgrading the OS and applications as service packs or upgrades are released; virus scanning and virus removal; backing up user data and software configuration data; and problem solving (both hardware and software). Problems may be caused by actions of the user, the software, or the computer system hardware. In most organizations, one or more administrators has the task of acquiring, installing, and maintaining the computer systems used by the organization.

While installing the OS and application software may be performed on the computer system prior to delivering the computer system to the user, other maintenance activities typically must be performed with the computer system at the user's location. In some cases, remote access may be used by the administrator to access the user's computer system remotely, so that the administrator need not physically travel to the user's location to operate on the computer system. However, some problems cannot be diagnosed remotely, particularly if the problem with the user's computer system is preventing the computer system from operating well enough to communicate on a network. When remote access is not possible, the administrator either travels to the user's location or temporarily moves the user's computer system to the administrator's location to correct

the problem. Costs associated with the traveling of the administrator and/or relocating the computer system may further increase TCO, particularly for larger organizations or geographically dispersed organizations.

5 One solution that has been used in the past is the "terminal server" solution. In this solution, the user's computer system is merely a "thin client" that includes software to connect to a central terminal server and to display an interface for the user. The user's applications are executed on the central terminal server, which is shared among many users. The central terminal server must be a relatively high powered (and thus expensive) 10 computer system, to provide acceptable processing power when all of the users are connected. Additionally, the central terminal server requires a fairly large amount of storage for user data. The terminal server solution centralizes most of the function in the organization, presumably near the administrators. However, the processing power of the user's computer system is frequently wasted, since the local processor in each computer 15 system is only used for display and communication with the central terminal server. Additionally, each terminal server can handle a maximum user load, and thus multiple terminal servers may be required. Furthermore, the terminal server solution is highly exposed to failures of the central terminal servers: when a given central terminal server fails, each user connected to the central terminal server experiences the failure, and is 20 down until the central terminal server can be brought back up.

In some cases, a network filesystem such as the Network Filesystem (NFS) is used and all applications on the user's computer system are configured to store configuration files and user data on networked storage. While centralization of user data is provided 25 this way (permitting, e.g., centralized backup of user data), many operating system configuration files cannot be stored in this fashion. For example, the Microsoft Windows™ operating system's registry is stored locally on the user's computer system. Accordingly, maintenance of the user's computer system still often involves accessing the user's computer system or relocation of the computer system. Furthermore, users may

choose to store some data locally on their computer systems, and such data is not backed up.

Portable computers, such as laptops, present additional challenges. Portable computers are often used away from the user's location, and thus the terminal server solution is not appropriate for the portable computer. Additionally, since the user physically carries the portable computer from location to location, the portable computer may be subject to physical loss or damage in addition to the typical problems experienced by fixed-location computer systems. While most data files may be backed-up from a portable computer system (either by the user or automatically, when it is connected to a network), many application configuration "tweaks" and other modifications may not be backed-up and thus may be lost when the portable computer system is lost or damaged.

SUMMARY OF THE INVENTION

15

In some embodiments, a system comprises at least one computer system, wherein the computer system is configured to execute a virtual machine corresponding to a user. The system further comprises a storage subsystem configured to store data representing the virtual machine and at least one file server. The file server is coupled to a network to which the computer system is configured to be coupled, and is also coupled to the storage subsystem. The file server is configured to provide the computer system with access to the data representing the virtual machine from the storage subsystem over the network.

In other embodiments, a computer accessible medium comprises a plurality of instructions which, when executed on a computer system responsive to a login of a user on the computer system, cause the computer system to execute a virtual machine corresponding to the user. The virtual machine is represented by data stored in a filesystem accessible to the computer system, at least intermittently. In still other embodiments, a computer system comprising the computer accessible medium and

execution hardware configured to execute the plurality of instructions is contemplated.

In some embodiments, a method is contemplated. Responsive to a login of a user on a computer system, a virtual machine is executed on the computer system. The virtual machine corresponds to the user. A filesystem is managed on a storage subsystem using at least one file server, wherein the storage subsystem stores data representing the virtual machine. At least intermittently, communication between the file server and the computer system over a network is performed to provide access to the data representing the virtual machine.

10

Another method is also contemplated in some embodiments. The method may be used, e.g., in a system comprising at least one computer system on which a user logs in during use, at least one file server coupled to a network to which the computer system is coupled at least intermittently, a storage subsystem storing data representing a virtual machine corresponding to a user, and a second computer system used by an administrator, wherein the file server manages a filesystem on the storage system and provides access to the data representing the virtual machine to the computer system. The method may comprise executing the virtual machine on the second computer system responsive to the user reporting a problem with the virtual machine; diagnosing the problem; and, if the problem is within the virtual machine, correcting the problem by modifying the data representing the virtual machine on the storage subsystem.

15

BRIEF DESCRIPTION OF THE DRAWINGS

20

The following detailed description makes reference to the accompanying drawings, which are now briefly described.

Fig. 1 is a block diagram of one embodiment of a system.

Fig. 2 is a block diagram of one embodiment of a client system that may be intermittently coupled to a network in the system of Fig. 1.

5 Fig. 3 is a block diagram of one embodiment of a client system that may be continuously coupled to a network in the system of Fig. 1.

Fig. 4 is a flowchart illustrating one embodiment of operating a client system that may be continuously coupled to the network in the system of Fig. 1.

10 Fig. 5 is a flowchart illustrating one embodiment of a boot sequence block from the flowchart of Fig. 4.

Fig. 6 is a flowchart illustrating one embodiment of a mount block from the flowchart of Fig. 4.

15 Fig. 7 is a flowchart illustrating one embodiment of a close block from the flowchart of Fig. 4.

20 Fig. 8 is a flowchart illustrating one embodiment of a method for diagnosing/correcting a user problem.

Fig. 9 is a flowchart illustrating one embodiment of operating a client system that may be intermittently coupled to the network in the system of Fig. 1.

25 Fig. 10 is a flowchart illustrating one embodiment of a determine volume state block from the flowchart of Fig. 9.

Fig. 11 is a table of volume states on the client and the file server for one embodiment of the flowchart shown in Fig. 9.

Fig. 12 is a block diagram illustrating one embodiment of a virtual machine configuration.

5 Fig. 13 is a block diagram illustrating one embodiment of a computer accessible medium.

While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will 10 herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

15

DETAILED DESCRIPTION OF EMBODIMENTS

Turning now to Fig. 1, a block diagram of one embodiment of a system comprising a plurality of computer systems assigned various roles is shown. Particularly, 20 a first set of client computer systems 10A-10N and a second set of client computer systems 12A-12M are shown. The client computer systems 10A-10N and 12A-12M (more briefly, "client systems") are computer systems assigned to users for their use. Additionally, a set of administrator computer systems (admin systems) 14A-14P are shown in Fig. 1. The admin systems 14A-14P are assigned to administrators for their use. 25 A set of file servers 16A-16Q, a maintenance server 18, and a provisioner server 20 are also shown in Fig. 1. A storage subsystem 22 is also shown. The client systems 10A-10N and 12A-12M, the admin systems 14A-14P, the file servers 16A-16Q, the maintenance server 18 and the provisioner server 20 are connected (or at least capable of being connected) to a network 24. The file servers 16A-16Q and the maintenance server

18 are coupled to the storage subsystem 22.

In the system of Fig. 1, each user may have at least one corresponding virtual machine dedicated to that user. In some embodiments, there may be a one-to-one correspondence between users and virtual machines. That is, when the user logs in to a client system 10A-10N or 12A-12M, the virtual machine that corresponds to that user may be started. In other embodiments, a given user may have more than one virtual machine. The user may log in to a client system 10A-10N or 12A-12M and select a virtual machine to be started. The user's virtual machine may comprise the machine state associated with that user, which may include the user's OS, application software, application and OS configuration data, and user data. Data representing each user's virtual machine may be stored on the storage subsystem 22. For example, in Fig. 1, the storage subsystem 22 may store data representing a set of virtual machines including VM1-VMn (reference numerals 28a-28n). The client systems 10A-10N and 12A-12M may be configured to execute virtual machines in response to the corresponding user's log in to a client system 10A-10N or 12A-12M.

The file servers 16A-16Q are computer systems that are configured to manage a filesystem on the storage system 22, and thus to provide user access to the virtual machine data on the storage subsystem 22. At least one file server 16A-16Q may be provided. Each file server 16A-16Q may include filesystem software (e.g. filesystem software 30 on the file server 16A) configured to manage the filesystem. In the illustrated embodiment, a plurality of file servers 16A-16Q are provided in a cluster 32 to provide high availability of the filesystem in the face of potential failures of file servers. That is, if a file server 16A-16Q fails, the client systems 10A-10N and 12A-12M that were using the failing file server 16A-16Q to access virtual machine data on the storage subsystem 22 may "fail over" to another file server 16A-16Q to continue access to the virtual machine data. In clustered embodiments, the file servers 16A-16Q may include cluster software (e.g. cluster software 34 in the file server 16A).

As mentioned above, the virtual machine data (sometimes referred to as the virtual machine image) contains the machine state associated with the user. For example, in the illustrated embodiment, the virtual machine data for each virtual machine may

5 include a configuration file (reference numeral 36) and one or more virtual disk files (reference numeral 38). The configuration file may store configuration information for the virtual machine, which may include the number of virtual disks, the size of the virtual disks, and any other information. The virtual disk files 38 may include the data stored on the virtual disks in the user's virtual machine. Thus, the virtual disk files 38 may include

10 some or all of the user's software (OS and applications), application and OS configuration files, etc. Additionally, the virtual disk files include some or all of the user's data. Each virtual disk may appear, to the user, as a storage device within the user's virtual machine.

15 While a configuration file 36 and a virtual disk file 38 for each virtual disk are included in the virtual machine data in the present embodiment, any organization of the virtual machine data may be used in other embodiments.

Since the virtual machine state is stored on the storage subsystem 22, maintenance of the users' virtual machines may be performed on the virtual machine data on the storage subsystem 22. For example, the maintenance server 18 may be included to

20 periodically perform various maintenance actions. As used herein, a "maintenance action" may comprise any activities which may be performed periodically on each user's machine (virtual machine, in this case) to help ensure the correct operation of the machines and/or the recovery of user state in the event of a failure. For example, maintenance actions may include, in some embodiments, one or more of: backup, virus

25 scan and virus removal, scan for corruption of data in the virtual disks (e.g. corrupt sectors, blocks, etc.), application of software updates/upgrades, etc. Performing the maintenance actions on the virtual machine data in the storage subsystem 22 may eliminate having to contact/connect to each client system 10A-10N and 12A-12M to perform the maintenance actions.

Problem diagnosis in the system of Fig. 1 may be performed by executing the virtual machine that exhibits the problem on an administrator's system 14A-14P. Once the problem is located and corrected, the administrator may stop the virtual machine and 5 the user may restart the corrected virtual machine from the storage subsystem 22. Traveling by the administrator to the user's location and/or relocating the user's computer system to the administrator's location may be eliminated, in some embodiments.

Since the client systems 10A-10N and 12A-12M execute the user's virtual 10 machines, including executing the user's software (OS and applications) within the user's virtual machine, the processing power of the client systems 10A-10N and 12A-12M may be more fully utilized, in some embodiments. The file servers 16A-16Q provide file service to the client systems 10A-10N and 12A-12M, but the execution of 15 applications/user OS is on the client systems 10A-10N and 12A-12M. Thus, the file servers 16A-16Q need only be provided with enough processing power to provide the file service, in some embodiments. Additionally, each user's virtual machine is independent of other virtual machines. If a given user's virtual machine is corrupted or otherwise cannot be executed, other users may not be impacted.

20 The client systems 12A-12M may be configured for essentially continuous connection to the network 24 during use. That is, once the client systems 12A-12M are booted and operating, the client systems 12A-12M are configured to maintain a connection to the network 24. The network 24 may be unavailable at times, but in general the client systems 12A-12M may have a network connection while in operation. 25 The client systems 12A-12M may be referred to more succinctly as "always-connected" systems.

A client system 12A-12M may access the virtual machine data for a given user on the storage subsystem 22 responsive to that user logging in to the client system 12A-12M.

Updates to the user's virtual machine (e.g. write operations to the user's virtual disks) may be written to the virtual machine data on the storage system 22. Reads of virtual machine data may be reads from the storage subsystem 22. In some embodiments, the client system 12A-12M may cache virtual machine data to offload the file servers 16A-16Q.

- 5 Since the client systems 12A-12M access a given virtual machine responsive to a log in by the corresponding user, any user may log in to any client system 12A-12M and may have access to the user's virtual machine on the client system 12A-12M.

The client systems 10A-10N may be intermittently connected to the network 24 (illustrated by the dotted lines 26A-26N in Fig. 1). For example, the client systems 10A-10N may include portable computer systems such as laptops that are sometimes operated without a network connection, as well as computer systems that are not portable but which have an intermittent network connection (e.g. a dialup modem connection). The client system 10A-10N may still execute a user's virtual machine. However, rather than relying on access to the storage subsystem 22, the client systems 10A-10N may be configured to replicate virtual machine data. For example, file replication or volume replication may be used. A copy of the user's virtual machine may be made to the client system 10A-10N, and the virtual machine may be executed from the local copy on the client system 10A-10N. Modifications to the virtual machine data may be replicated by the client system 10A-10N to the virtual machine data on the storage subsystem 22 when a network connection is available. Similarly, if a maintenance action is performed on the virtual machine on the storage subsystem 22 or a problem is corrected by an administrator and the corrected virtual machine is on the storage subsystem 22, replication from the storage subsystem 22 to the client system 10A-10N may be performed to provide the updated/corrected virtual machine to the client system 10A-10N.

In addition to using replication for intermittently-connected client systems, replication may also be used for client systems 10A-10N that may have an essentially continuous network connection that is not "fast enough" (e.g. not high enough bandwidth

or with a latency that is too high) to permit acceptable performance if the virtual machine is operated in the manner described above for the client systems 12A-12M. Since the replication may occur in the background, user operation of the virtual machine may not be slowed by the network connection speed even though the user's updates to the virtual machine data are being replicated. In some embodiments, all client systems may use replication rather than executing the virtual machine from the storage subsystem 22. User log in to any client system may still be supported. When the user logs in to a different client system than the client system that the user logged in to the previous time, a replication of the user's virtual machine data to the new client system may occur.

10

- In some embodiments, the provisioner server 20 may be provided. The provisioner server 20 may be configured to provision client systems 10A-10N and 12A-12N with the resources to execute virtual machines. More particularly, if an administrator corrects a problem in a user's virtual machine, the administrator may also choose to reprovision the client system that was executing the user's virtual machine (e.g. if the problem indicates that the software or configuration of the client system may contribute to the problem). Additionally, in some cases, the client system hardware may be replaced (e.g. due to a problem identified by a user, to retire old hardware, etc.). The new client system may be provisioned on its initial boot by the provisioner server 20.
- 20 The provisioner server 20 may include provisioner software that is designed to provision computer systems with resources (which may include one or more of OS software, application software, and/or various configuration information such as Internet Protocol (IP) address, etc.). Any provisioner software may be used in various embodiments. For example, the OpForce™ products provided by VERITAS™ Software Corporation (Mountain View, CA) may be used. Other examples may include provisioning products from Altiris, Inc. (Lindon, UT), netboot, provisioning products from Opsware, Inc. (Sunnyvale, CA), PowerCockPit from Mountain View Data, Inc. (Mountain View, CA), and Automated Deployment Services from Microsoft Corp. (Redmond, WA), etc. In some embodiments, the provisioner server 20 may be coupled to the storage subsystem 22

as well. For example, the provisioner server 20 may store resource images to be installed on various client systems on the storage subsystem 22.

- The storage subsystem 22 may comprise any number and type of storage device.
- 5 For example, the storage subsystem 22 may comprise storage devices in a storage area network (SAN) configuration, and the storage subsystem 22 may include SAN switches and other devices for providing connectivity to the SAN storage devices. Other embodiments may comprise network attached storage (NAS) and/or storage devices coupled to a peripheral interface such as small computer systems interface (SCSI), Fibre
- 10 channel, integrated drive electronics (IDE) interface, other storage controllers, etc. Any combination of storage devices coupled in any desired fashion may be used.

- The file servers 16A-16Q may comprise computer systems configured to provide file service for the storage subsystem 22. Any filesystem capable of service over a
- 15 network may be used. In one particular implementation, for example, NFS may be used. NFS may be stateless on the file server (that is, no file state may be maintained on the NFS servers). Thus, fail over from one NFS server to another may, in some cases, be relatively straightforward and relatively rapid, since no file state needs to be recovered from the failing NFS server. In other implementations, other filesystems may be used
- 20 (e.g. Andrew filesystem (AFS), common Internet filesystem (CIFS), etc.). The filesystem software 30 on the file server 16A (and similar software on other file servers) may implement the filesystem.

- In one embodiment, each user may have a directory in the filesystem on the
- 25 storage subsystem 22, in which the virtual machine data representing the user's virtual machine is stored. The user's directory may be mounted by the client system that the user logs in to. In the case of client systems 10A-10N, the user's directory may be mounted if the client system 10A-10N is connected to the network 24.

As mentioned above, in some embodiments, the file servers 16A-16Q may be clustered for high availability. A cluster may comprise two or more computer systems that permit fail over from one computer system to another in a fashion that is essentially transparent to client systems coupled to the cluster. For example, a given file server 16A-
5 16Q may provide file service to a given client system 10A-10N or 12A-12M. If the given file server 16A-16Q experiences a failure, the file service for the given client system may be failed over to another file server 16A-16Q without the given client system having to reconnect to the file server. While filesystem performance may be temporarily affected, the given client system may otherwise be unaware of the fail over and may continue to
10 receive file service from the failed-to file server. In embodiments that implement a cluster, the cluster software 34 may be included on each file server 16A-16Q. Any cluster software may be used. For example, VERITAS Cluster Server™ (VCS) available from VERITAS Software Corporation may be used in some embodiments.

15 Generally, a virtual machine may comprise any combination of software, one or more data structures in memory, and/or one or more files on one or more storage devices. The virtual machine represents the software and hardware used by a user to which the virtual machine is assigned. For example, the virtual machine may include the user's data, applications, and OS software. The configuration file of the virtual machine may
20 specify a virtual CPU, the virtual disks, and other virtual input/output devices (e.g. virtual network interface cards), the amount of memory allocated to the virtual machine, etc.

Different virtual machines which execute on the same computer system may differ. For example, the OS included in each virtual machine may differ. Different
25 virtual machines may employ different versions of the same OS (e.g. Microsoft Windows NT with different service packs installed), different versions of the same OS family (e.g. Microsoft Windows NT and Microsoft Windows 2000), or different OSs (e.g. Microsoft Windows NT, Linux, Sun Solaris, other Unix flavors, etc.). Furthermore, the OS in each virtual machine may differ from the OS installed on the client computer systems.

The network 24 may comprise any network technology in various embodiments. The network 24 may be a local area network, metropolitan area network, wide area network, intranet network, Internet network, wireless network, or any other type of 5 network or combinations of the above networks. Any network media may be used. For example, the network 24 may comprise an Ethernet network. Alternatively, the network may comprise a token ring network, etc.

In various embodiments, the number of each type of computer system shown in 10 Fig. 1 may vary. That is, any number of client systems 10A-10N may be included; any number of client systems 12A-12M may be included; any number of admin systems 14A-14P may be included; any number of provisioner servers 20 may be included; any number of file servers 16A-16Q may be included; and any number of maintenance servers 18 may be included.

15 Turning now to Fig. 2, a block diagram of one embodiment of the client system 10A is shown. Any intermittently-connected client system or slow-connected client system may be similar. In the illustrated embodiment, the client system 10A includes hardware 40 (which further includes a storage device 42), filesystem client software 44, replication software 46, a client OS 48, a virtual machine (VM) kernel 50, and a login 20 script 52. The storage device 42 includes a replicated volume 54, on which the data representing the user's VM 56 is stored. In other embodiments, the replicated volume 54 is optional (e.g. if file replication is used).

25 The hardware 40 generally comprises the hardware included in the client system 10A. The hardware 40 may comprise one or more processors configured to execute the instructions comprising the software on the client system 10A, memory for use by the software on the client system 10A, and peripheral devices (e.g. storage devices such as device 42, network connection hardware, user interface devices, etc.).

The filesystem client software 44 may comprise software for interacting with the filesystem software on the file servers 16A-16Q. In some embodiments, the filesystem client software 44 may comprise a driver for the client OS 48.

5

The replication software 46 comprises the software used to replicate the virtual machine data representing the user's virtual machine to/from the storage subsystem 22. In the illustrated embodiment, the replication software 46 may perform volume replication. A volume on the storage device 42 (or on multiple storage devices 42) may be defined 10 and may be identified to the replication software 46 as a volume to be replicated to the storage subsystem 22. The files representing the user's VM 56 may be stored in the replicated volume. The replication software 46 may be configured to log updates to the volume and to replicate the updates to the storage subsystem 22 when a network connection is available. Thus, the files representing the user's VM 56 may be replicated 15 to the storage subsystem 22. Any volume replication software may be used. For example, in some embodiments, the VERITAS Volume ReplicatorTM available from VERITAS Software Corporation may be used. Other examples may include Symmetrix Remote Data Facility from EMC Corporation. In other embodiments, file replication software may be used. File replication software replicates on a file basis. That is, when a 20 file is updated, the change to the file is replicated. In some implementations, the entire file may be replicated. In other embodiments, only the change to the file may be replicated. Any file replication technology may be used (e.g. incremental, bulk, etc.). The files representing the user's VM 56 may be identified to the file replication software to be replicated to the storage subsystem 22. Any file replication software may be used. 25 For example, VERITAS Storage ReplicatorTM available from VERITAS Software Corporation may be used. In other embodiments, other file replication software may be used (e.g. DoubleTake from NSI Software). Other embodiments may use the EMC Replication Manager from EMC Corporation.

It is noted that, while Fig. 2 illustrates the user's virtual machine data as residing in a replicated volume, multiple replicated volumes may be used to store different portions of the user's virtual machine data in some embodiments. For example, in some embodiments, one or more virtual disks may be stored on different replicated volumes.

5

The client OS 48 is the OS executed by the client system 10A. Any OS may be used (e.g. any version of Microsoft Windows, Sun Solaris, Linux, other Unix varieties, etc.). Particularly, the client OS 48 may differ from the user's OS included in the user's virtual machine.

10

The VM kernel 50 may comprise the software that schedules virtual machines for execution on the underlying hardware 40 and handles traps from the virtual machines (e.g. when virtual hardware is accessed, instruction execution exceptions, etc.). In one embodiment, the VM kernel 50 may comprise the GSX product available from VMWare, Inc. (Palo Alto, CA), now owned by EMC Corporation. In another embodiment, the VM kernel may comprise the Virtual PC product available from Connectix Corporation (San Mateo, CA), now owned by Microsoft Corp. The VM kernel may include a Java Virtual Machine, in some embodiments. In other embodiments, the VM kernel may include virtual machine technologies for the Linux platform such as user mode Linux (UML) virtual machines or plex86. It is noted that the VM kernel may also be referred to as a virtual machine monitor (VMM) such as the VMMs used on mainframe computer systems such as those available from International Business Machines Corporation. In the illustrated embodiment, the VM kernel may execute on the client OS 48. In other embodiments, the VM kernel may execute directly on the underlying hardware (i.e. without an underlying operating system). For example, the ESX product available from VMWare, Inc. may be used.

The login script 52 may comprise software that executes to present a log in interface to a user, to authenticate the user, and to start execution of the user's virtual

machine. The login script 52 may be written in a script language (e.g. perl, korn, c shell, etc.), or may comprise compiled software instructions, in various embodiments.

It is noted that the client system 10A may include a computer accessible medium 5 storing the filesystem client software 44, the replication software 46, the client OS 48, the VM kernel 50, and the login script 52 (not shown in Fig. 2).

Turning now to Fig. 3, a block diagram of one embodiment of the client system 12A in more detail is shown. In the illustrated embodiment, the client system 12A 10 includes the hardware 40, the filesystem client software 44, the client OS 48, the VM kernel 50, and the login script 52.

In some embodiments of the client system 12A, the filesystem client software 44 15 may implement a caching mode. In the caching mode, recently used and updated file blocks may be cached on the client system, and the cache may be checked for a hit whenever a block is accessed by the client system 12A. The caching may be controlled on a file-by-file basis, in some embodiments (e.g. caching may be enabled for files containing virtual machine data and disabled for other files). For updates, the block may be added to the cache after the update is complete in the storage subsystem 22. 20 Additional details regarding one embodiment of the caching mode are provided below.

It is noted that the client system 12A may include a computer accessible medium 25 storing the filesystem client software 44, the client OS 48, the VM kernel 50, and the login script 52 (not shown in Fig. 3).

Turning now to Fig. 4, a flowchart illustrating operation of one embodiment of a client system 12A-12M (an "always-connected client") is shown. Blocks that are implemented in software may represent a plurality of instructions comprising the software which, when executed on the always-connected client, perform the function(s) described

for the blocks.

The always-connected client may proceed through a boot sequence when initially powered-on (block 60). Generally, the term "boot" refers to any set of events, hardware and/or software, used to initialize a computer system and load the operating system (e.g. the client OS 48) for execution. In some embodiments, the boot sequence may include communicating with the provisioner server 20 to determine if the always-connected client is to be provisioned. An exemplary embodiment is shown in Fig. 5 and described in more detail below.

10

Once the always-connected client has booted, the login script 52 may be executed. The login script 52 may present an interface for a user to log in to the always-connected client. A user logs in and is authenticated (block 62). Any mechanism for authenticating a user may be used. For example, the user's password may be checked against a locally-stored password file such as /etc/passwd. Alternatively, Microsoft's Active Directory authentication may be used. In other alternatives, lightweight directory access protocol (LDAP) authentication or Network Information Service (NIS, also known as yellow pages or directory service) authentication may be used.

20

Once the user is authenticated, the login script 52 may mount the user's directory or directories from the filesystem provided by the file servers 16A-16Q (block 64). One embodiment of block 64 is shown in Fig. 6 and described in more detail below. If the user's directory is successfully mounted, the login script may start execution of the user's virtual machine (block 66). Execution of the user's virtual machine, accessing and updating the files in the user's directory on the storage subsystem 22, continues until the user exits (decision block 68). Once the user decides to exit, the always-connected client closes the user's virtual machine (block 70). One embodiment of block 70 is shown in Fig. 7 and described in more detail below. If the always-connected client is being shutdown (decision block 72, "yes" leg), the flowchart exits. If the always-connected

client is not being shutdown (decision block 72, "no" leg), the login script 52 may be executed again to present the login interface to the next user.

Fig. 5 is a flowchart illustrating one embodiment of the boot sequence (block 60 in Fig. 4) used by one embodiment of an always-connected client. Blocks that are implemented in software may represent a plurality of instructions comprising the software which, when executed on the always-connected client, perform the function(s) described for the blocks. For example, the blocks in Fig. 5 may be implemented in firmware on the always-connected client, such a basic input/output system (BIOS) code.

10

In the embodiment of Fig. 5, the boot sequence may include transmitting a remote boot request (block 80). For example, personal computer (PC) computer systems may implement a Pre Execution Environment (PXE) boot request. A PXE boot request may comprise a specially formatted packet transmitted on the network 24 that identifies the booting computer system (e.g. by the media access controller (MAC) address in the network interface controller in the booting computer system) and indicates that a boot is occurring. The booting computer system may check for a response on the network, and may timeout after a specified interval if no response is received. If a timeout occurs, boot may continue with other bootable devices (e.g. fixed or removable disk drives in the booting computer system or coupled thereto). In other embodiments, other types of remote boot protocols may be used. For example, computer systems available from Sun Microsystems, Inc. (Santa Clara, CA) may support a net boot protocol. As used herein, a "remote boot request" may comprise any communication from a booting computer system, transmitted external to the computer system in order to receive boot code from a remote system or device. A response to the remote boot request may include the boot code, or may indicate where the boot code may be located (within the booting computer system or external to the booting computer system).

In the present embodiment, the provisioner server 20 may respond to the remote

boot request if the booting always-connected client is to be provisioned. For example, the provisioner server 20 may record an indication for each computer system that is to be provisioned. In one particular embodiment, the provisioner server 20 may maintain a database of the computer systems in the system of Fig. 1, and other systems (e.g. admin systems 14A-14P) may update the database to indicate which computer systems are to be provisioned. If the provisioner server 20 responds to the remote boot request (decision block 82, "yes" leg), the always-connected client receives the provisioned resources from the provisioner server 20 (block 84), installs the provisioned resources, and reboots (returning to block 80). In one implementation, the provisioned resources may comprise the resources shown in Fig. 3.

If the provisioner server 20 does not answer the remote boot request (decision block 82, "no" leg), the always-connected client may determine if the timeout for the remote boot request has expired (decision block 86). If the timeout has not expired (decision block 86, "no" leg), the always-connected client may continue to await a response from the provisioner. If the timeout has expired (decision block 86, "yes" leg), the always-connected client may boot from the local disk (block 88).

Fig. 6 is a flowchart illustrating one embodiment of mounting the user's directory (block 64 in Fig. 4) that may be used by one embodiment of an always-connected client. Blocks that are implemented in software may represent a plurality of instructions comprising the software which, when executed on the always-connected client, perform the function(s) described for the blocks. For example, the blocks in Fig. 6 may be implemented in the login script 52.

25

The login script 52 may determine which file server 16A-16Q to use to mount the users directory (block 90). Any mechanism may be used. For example, a local configuration script may be executed, NIS may be used, LDAP may be used, Active Directory may be used, or a lookup in a central database may be used. The login script 52

may mount the user's directory (block 92). The login script 52 may check the mount to ensure that the virtual machine may be executed from the mount. For example, the login script 52 may check if the mount is read-only. If the mount is read-only, then updates may not be propagated to the storage subsystem 22. Accordingly, if the mount is read-only (decision block 94, "yes" leg), the login script 52 may inform the user, unmount the user's directory, and log the user off (block 96). The login script 52 may then return to the login block 62 in Fig. 4 (block 98). Another check may determine if the mount is locked (decision block 100). For example, a lock file may be used in the present embodiment. If a lock file exists in the user's directory, then another system may have the user's directory mounted (e.g. if the user has not logged off the other system, or if there is a problem with the other system or the user's virtual machine that has prevented log off on the other system). Other embodiments may use other locking schemes. If the mount is locked (decision block 100, "yes" leg), the login script 52 may optionally identify which computer system 10A-10A or 12A-12M has the lock and may give the user the option to "break" the lock and continue operation (decision block 101). Other embodiments may not provide the option to break the lock. Permitting the user to break the lock may be used to recover from an error or other problem on another computer system that the user was previously using, although a loss of data may occur. If the user does not choose to break the lock (decision block 101, "no" leg) (or if the option to break the lock is not supported), the login script 52 may inform the user, unmount the user's directory, and log the user off (block 96). The login script 52 may then return to the login block 62 in Fig. 4 (block 98). If the user chooses to break the lock (decision block 101, "yes" leg), operation may continue at block 102. In some cases, breaking the lock may also include other cleanup, e.g. changing the identity of the last always-connected client, before continuing operation.

If the mount is not locked, the login script 52 may lock the user's directory atomically (block 102). For example, if a lock file is used, the login script 52 may create the lock file in the user's directory. In some embodiments, the contents of the lock file

may be the identity of the always-connected client that locked the directory. In such embodiments, an administrator may be able to determine which always-connected client is executing a user's virtual machine.

5 In some embodiments, a virtual disk may be dedicated for swapping or paging of virtual memory by the user's OS. As used herein, the term "swap" refers to the OS's movement of virtual memory information between memory and external storage, and may refer to the external storage that contains the virtual memory information in some cases. For example, if the user's OS is Windows, the swap file may be allocated to the dedicated
10 virtual disk. For various Unix implementations, the swap partition may be allocated to the dedicated virtual disk. Dedicating a virtual disk to swapping may permit the swap data to be stored only locally on the always-connected client. The dedicated virtual disk on the storage subsystem 22 may comprise configuration information describing the virtual disk. For example, the configuration information may include its size, and
15 indication that it is the dedicated swap virtual disk, etc. The dedicated virtual disk may be copied to the local storage in the always-connected client, and updates to the dedicated virtual disk on the local storage may not be propagated to the storage subsystem 22. Thus, paging may not create traffic to the filesystem 30 and the storage subsystem 22.
The swap data is not required after the virtual machine is closed, and so reflecting swap
20 updates to the storage subsystem 22 may not be required. In such embodiments, the login script 52 may copy the dedicated swap virtual disk to the local disk (block 104).

In some embodiments, the last always-connected client to execute the user's virtual machine may be recorded. For example, a last-used file may be created in the
25 user's directory, and the contents of the file may identify the last always-connected client to execute the virtual machine. The login script may check the last-used file to determine if the current always-connected client was the last to execute the virtual machine. If not (decision block 106, "no" leg), the local cache of the always-connected client may be cleared (block 108). If so (decision block 106, "yes" leg), the local cache of the always-

connected client may contain cache data that may be used by the executing virtual machine. Thus, the local cache may not be cleared. In either case, if caching mode is implemented, the local cache may be enabled (block 110).

5 Turning next to Fig. 7, a flowchart illustrating one embodiment of closing the user's virtual machine (block 70 in Fig. 4) that may be used by one embodiment of an always-connected client is shown. Blocks that are implemented in software may represent a plurality of instructions comprising the software which, when executed on the always-connected client, perform the function(s) described for the blocks.

10

The always-connected client may disable the local file cache (block 120). Additionally, the always-connected client may record itself as the last always-connected client to execute the user's virtual machine (block 122) and may unlock the user's directory (block 124). For example, in embodiments that use the lock file and the last-used file, the always-connected client may rename the lock file as the last-used file (thereby unlocking the user's directory and recording itself as the last-used client). The always-connected client may unmount the user's directory (block 126). The always-connected client may delete the local swap file/partition (block 128) to maintain security, since some of the user's information may be in the local swap file/partition. The always-connected client may log off the user (block 130).

25 While some embodiments described above use a lock file to implement locking, other embodiments may implement locking in other fashions. For example, a database may be maintained of directories and whether or not they are locked. Alternatively, a central locking resource may be implemented, and client systems may communicate with the central locking resource to acquire and release locks. As another example, the file servers 16A-16Q may implement the locking. Similarly, while a last-used file is used in some embodiments above to record the last client system to execute a virtual machine, other embodiments may use other methods (including methods similar to the above-

mentioned examples of locking). Locking and/or recording the last client system to execute a virtual machine may be optional, and may not be implemented in other embodiments.

5 Turning next to Fig. 8, a flowchart illustrating one embodiment of a method for diagnosing and correcting a user problem is shown. Blocks that are implemented in software may represent a plurality of instructions comprising the software which, when executed, perform the function(s) described for the blocks.

10 Generally, a user problem may encompass any problem which prevents a user from executing his/her virtual machine, prevents him/her from successfully using applications installed in the virtual machine, or impacts the performance of the virtual machine. The problem may be in the user's software (within the virtual machine), the software on the client that the user logged into, the client hardware, or a combination
15 thereof.

The user may exit the virtual machine and/or power off the client system (block 140). If the user is able to exit the virtual machine properly (e.g. the virtual machine and/or client system is in operable enough state to exit), the user need not power off the
20 client system unless instructed to do so by the administrator. On the other hand, if the user is not able to exit the virtual machine, the user may power off the client system. The user may notify the administrator of the problem (block 142).

The administrator, on one of the admin systems 14A-14P, starts the user's virtual
25 machine (block 144). In some embodiments, the administrator may need to take steps to force the user's virtual machine to start (e.g. clearing locks that were not cleared due to a failed exit of the virtual machine, etc.). The administrator diagnoses any problems that may exist in the user's virtual machine and corrects the problems (block 146). The administrator exits the user's virtual machine, and thus the corrected virtual machine is

stored on the storage subsystem 22.

Optionally, if the problem may involve software installed on the client system, the administrator may inform the provisioner server 20 that the client system is to be reprovisioned (block 148). For example, if the provisioner server 20 maintains a database of all the computer systems in the system of Fig. 1, the administrator may update the database to indicate that the client system is to be reprovisioned. In another example, the administrator may reprovision the client system as a default, to assure that any contributing factors on the client system are eliminated. The administrator then informs the user that the problem is fixed (block 150). If the system is to be reprovisioned, the administrator may request that the user reboot the client system. The user powers up the client system (block 152) and attempts to log in and start the virtual machine. If the virtual machine starts properly (decision block 154, "yes" leg), the problem may be viewed as fixed. If the virtual machine does not start properly (decision block 154, "no" leg), the client system hardware may be defective (block 156). The hardware may be called in for repair, or simply replaced.

Turning now to Fig. 9, a flowchart illustrating operation of one embodiment of a client system 10A-10N (an "intermittent client") is shown. Blocks that are implemented in software may represent a plurality of instructions comprising the software which, when executed on the intermittent client, perform the function(s) described for the blocks. The embodiment illustrated in Fig. 9 describes volume replication for replicating changes to the user's virtual machine data. Other embodiments may employ file replication.

The intermittent client may boot locally (block 160). That is, the intermittent client may not issue a remote boot request as part of its boot sequence. Since the intermittent client may not always be connected to the network, the remote boot request may frequently timeout, delaying the boot unnecessarily. Alternatively, the user may be given the option to transmit a remote boot request, and the intermittent client may boot

locally or may use a boot sequence similar to that shown in Fig. 5 based on the user's response.

The login script 52 on the intermittent client may execute and present a login interface for the user. The user may login and be authenticated (block 162). Any authentication method may be used, although the method may default to a local method if no network connection is available. If a network connection is available, the network connection may be established (including any user input that may be required to establish the network connection).

10

The login script 52, in concert with the replication software 46, may determine the volume state (block 164). As part of determining the volume state, and dependent on whether or not a network connection is available, login script 52/replication software 46 may prompt the user to determine whether or not to continue. A more detailed discussion 15 of one embodiment of determining the volume state is provided below with respect to Fig. 10. Assuming that the determination of the volume state permits continued execution (see, e.g., Fig. 10), the login script 52 may start the user's virtual machine (block 166). The virtual machine may be started from the virtual machine data on the intermittent client (e.g. the data in the replicated volume(s), in the volume replication 20 embodiments). It is noted that, while the present embodiment is described with respect to volume replication, other embodiments may implement file replication. A similar determination of the state of the files may be performed in such embodiments.

The replication software 46 may generally log changes to the replicated volume(s) 25 (that is, to the virtual machine data) (block 168) as the virtual machine continues executing on the intermittent client. Additionally, if a network connection is available, the replication software 46 may replicate the changes to the storage subsystem 22. The process of logging changes and replicating changes may continue as the virtual machine execution continues (decision block 170, "no" leg). If the user exits the virtual machine

(decision block 170, "yes" leg), replication may continue until the volume replication completes. That is, replication may continue until the virtual machine data in the storage subsystem 22 is up to date with the virtual machine data on the intermittent system. If the volume replication is complete (decision block 172, "yes" leg), the replication software 46 may change the volume states on the intermittent system and the storage subsystem 22 to idle (block 174). The intermittent system may wait for the next user login (block 162). If the volume replication is not complete (decision block 172, "no" leg), the replication software 46 may continue replication (block 176). At any time while replication continues after the user exits the virtual machine, a user may log in to the intermittent system. Replication may continue in the background after the user logs in.

Turning now to Fig. 10, a flowchart is shown illustrating one embodiment of determining the volume state (block 164). Blocks that are implemented in software may represent a plurality of instructions comprising the software which, when executed on the intermittent client, perform the function(s) described for the blocks. Volume state may be determined relative to the volume replication. There may be at least four states for a replicated volume: primary connected, primary disconnected, secondary, and idle. The primary connected state or the primary disconnected state are states of the replicated volume that is actively updated by the user. For example, during normal operation, the volume on the intermittent system may be in either the primary connected state or the primary disconnected state. The primary connected state is used if a network connection to a secondary volume exists, and the primary disconnected state is used if such a network connection does not exist. A secondary state applies to a volume receiving replication from the primary volume. For example, during normal operation, the volume on the storage subsystem 22 is the secondary volume. The idle state applies to a volume in which replication is not currently underway.

The replication software 46 may determine if a network connection exists (decision block 180). If a network connection exists (decision block 180, "yes" leg), the

replication software 46 may determine the remote and local volume states and determine if the states are appropriate for starting operation (decision block 182). The local volume state may be the state of the volume on the intermittent client. The remote volume state may be the state of the volume on the storage subsystem 22. Fig. 11 is a table illustrating 5 possible volume states and the result for one embodiment of volume replication. Fig. 11 is described in more detail below.

If the remote and local volume states are appropriate for starting operation (decision block 182, "yes" leg), the replication software 46 may generally begin 10 operation. However, if the remote volume state is primary and the local volume state is secondary (decision block 184, "yes" leg), the replication software 46 is in the process of replicating from the storage subsystem 22 to the intermittent client. The replication software 46 may wait for the replication to complete (block 186). In either case, the replication software 46 may change the local volume state to primary connected, the 15 remote volume state to secondary of the intermittent client (block 188).

On the other hand, if the remote and local volume states are not appropriate for starting operation, such as the error cases in Fig. 11, the replication software 46 may wait for the error to be corrected before continuing (decision block 182, "no" leg and block 20 190). In some embodiments, the error may be corrected manually, such as by an administrator. In some cases, the definition of the virtual disks in the user's virtual machine may simplify correction of the error.

If a network connection is not available (decision block 180, "no" leg), the remote 25 volume state may not be determined. The replication software 46 may inform the user of the local volume state and request input (block 192). The replication software 46 may also indicate the age of the local volume state. If the user elects not to continue (decision block 194, "no" leg), the login script 52 may log the user off (block 196). Execution may return to block 162 to await user login (block 198). If the user elects to continue

(decision block 194, "yes" leg), the replication software 46 may change the state to primary disconnected (block 200).

The user may consider a variety of factors in deciding whether or not to continue
5 if the network connection is not available. If the local volume state is primary and the state is recent, the user may conclude that the likelihood of more recent updates on the storage subsystem 22 is low and the user may elect to continue. If the local volume state is secondary and is recent, the user may conclude that replication from the storage subsystem 22 to the intermittent client was in progress recently and not completed, and
10 thus the user may choose not to continue. In other embodiments, the login script 52/replication software 46 may default whether or not to start the virtual machine using the above rules or other rules, and may not request input from the user.

It is noted that, in embodiments in which a client system 10A-10N has a network
15 connection but it is not fast enough to provide acceptable performance as an always-connected client, a network connection is available and blocks 180, 192, 194, 196, 198, and 200 may be eliminated from the flowchart of Fig. 10.

Turning next to Fig. 11, a table of volume states and actions for one embodiment
20 of the replication software 46 is shown.

The first row of the table in Fig. 11 illustrates the volume states for normal operation when the intermittent system is connected to the network 24. The local volume state is primary connected, and the remote volume state is secondary to the intermittent client system. In this case, the volume replication software 46 may continue replication from the primary volume to the secondary volume.
25

The second row of the table illustrates a first error case for the volume states, in which the primary volume state is primary connected and the remote volume state is not

secondary of the intermittent client system. This error case (denoted Error1 in the table) may occur if the remote volume is forced out of the secondary state. For example, the remote volume may be forced out of the secondary state if an administrator makes a correction to the user's virtual machine data to fix a problem. Alternatively, a

- 5 maintenance action may force the remote volume out of secondary state (e.g. a software update, a virus removal, etc.).

The third row of the table is another normal case in which the local volume state is primary disconnected and the secondary volume state is idle. In this case, the
10 replication software 46 may change the remote volume state to secondary of the intermittent client system, the local volume state to primary connected, and begin replication.

The fourth row of the table is a second error case (labeled Error2 in the table) in
15 which the local volume state is primary disconnected and the remote volume state is not idle. This error case may occur, for example, if the virtual machine is started on the intermittent client when the intermittent client is not connected to the network.

The fifth row of the table is the case of replicating from the storage subsystem 22
20 to the intermittent client. In this case, the local volume state is secondary and the remote volume state is primary of the intermittent client system. This case may occur, for example, if a change to the virtual machine data is made on the storage subsystem 22 and the volume states are set to replicate the update to the intermittent client (either manually or automatically as part of an update process in which the user exits the virtual machine,
25 the replication to the storage subsystem 22 is completed, and then the update is subsequently performed on the storage subsystem 22). This case may also be forced by an administrator after correcting a problem with the user's virtual machine, to cause the corrected virtual machine to be copied to the intermittent client.

The sixth row of the table is a third error case (labeled Error3 in the table). In this case, the local volume state is secondary and the remote volume state is not primary of the intermittent client. This case may occur, for example, if the intermittent client is absent from the network 24 for an extended period and the states are reset.

5

The seventh row of the table is the case in which both volume states are idle. This is the case after replication is completed, for example.

In one embodiment, the error cases may be corrected in one of three ways: (1) the
10 virtual machine data on the intermittent client may override the virtual machine data on the storage subsystem 22; (2) the virtual machine data on the storage subsystem 22 may override the virtual machine data on the intermittent client; or (3) a new virtual machine may be created and virtual machine data from either system may be used as desired. In some cases, the definition of the virtual machine may make simplify the decision.

15

For example, Fig. 12 is an exemplary virtual machine 210 including a configuration file 36 and four virtual disks 38 (labeled VD1 through VD4). VD1 may comprise the user's software, including the user's OS 212, other system software 214, and applications 216. VD2 may comprise user files 218 (e.g. files creates using the user's
20 applications). VD3 may comprise the swap file or swap partition 220, as described above. VD4 may comprise read-only shared data 222 that may be shared among multiple users.

Separating user data on VD2 (typically modified when the virtual machine is
25 executing on the intermittent client) from software on VD1 (typically modified by the maintenance server or an administrator on the storage subsystem 22) may simplify the decision of how to correct the error cases. For example, the storage subsystem 22 may override the intermittent client for VD1 and the intermittent client may override the storage system 22 for VD2. For volume replication, VD1 and VD2 may be on separate

volumes. Alternatively, file replication may be used and the VD1 file from the storage subsystem 22 and the VD2 file from the intermittent client may be selected.

Providing a virtual disk for read-only shared data (VD4) may provide a
5 convenient mechanism for sharing data among multiple users. For example, in some embodiments, the VD4 file may be a symbolic link included in any virtual machine for which shared access to the files on the VD4 virtual disk is desired. Only one copy of the files may be provided in such an embodiment. In other embodiments, the files may be included directly in the VD4 virtual disk in each virtual machine. In still other
10 embodiments, a virtual disk or disks may provide shared access to data and write access may also be permitted. For example, each virtual machine may make a copy of the shared data and may make the copy read/write. Alternatively, each virtual machine may create one or more copy-on-write (COW) files to store changes made to the shared virtual disk(s) in that virtual machine. The COW files may be part of each individual virtual
15 machine. While updates to the shared data may not be consistent across various virtual machines, such consistency may not be desired in some circumstances. For example, the shared data may comprise a common software installation used in multiple virtual machines, and the COW file(s) in each virtual machine may contain user customizations of the common installation.

20

Turning now to Fig. 13, a block diagram of a computer accessible medium 230 is shown. Generally speaking, a computer accessible medium may include any media accessible by a computer during use to provide instructions and/or data to the computer. For example, a computer accessible medium may include storage media such as magnetic or optical media, e.g., disk (fixed or removable), CD-ROM, or DVD-ROM, CD-R, CD-RW, DVD-R, DVD-RW, volatile or non-volatile memory media such as RAM (e.g. synchronous dynamic RAM (SDRAM), Rambus DRAM (RDRAM), static RAM (SRAM), etc.), ROM, Flash memory, non-volatile memory (e.g. Flash memory) accessible via a peripheral interface such as the Universal Serial Bus (USB) interface,

etc., as well as media accessible via transmission media or signals such as electrical, electromagnetic, or digital signals, conveyed via a communication medium such as a network and/or a wireless link. The computer accessible medium 230 in Fig. 13 may be encoded with one or more of the VM kernel 50, the login script 52, the filesystem client software 44, the replication software 46, the filesystem software 30, the cluster software 34, and/or the virtual machines 28a-28n and 210. Each of the VM kernel 50, the login script 52, the filesystem client software 44, the replication software 46, the filesystem software 30, and the cluster software 34 may comprise instructions which, when executed, implement the functionality described herein for the software. Generally, the computer accessible medium 230 may store any set of instructions which, when executed, implement a portion or all of the flowcharts shown in one or more of Figs. 4-10. In some embodiments, a computer accessible medium similar to the computer accessible medium 230 may be included in a client system 10A-10N (e.g. the embodiment of Fig. 2), a client system 12A-12M (e.g. the embodiment of Fig. 3), and/or a file server 16A-16Q.

15

Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.